arXiv:1703.10272v1 [cs.DC] 29 Mar 2017

Fast and Flexible Data Analytics with F^2

Robert Grandl, Arjun Singhvi and Aditya Akella
University of Wisconsin-Madison

***Abstract*—**Existing data analytics frameworks are intrinsically compute-centric in nature. Their computation structure is complex and determined early, and they take decisions that bind early to this structure. This impacts expressiveness, job performance, and cluster efficiency.

We present F^2, a new analytics framework that separates computation from data management, making the latter an equal first-class entity. We argue that this separation enables more flexibility in expressing analytics jobs and enables data driven optimizations. Furthermore, it enables a new kind of "tasks" with loose semantics that can multiplex their execution across different sets of data and multiple jobs.

1 Introduction

Data analytics drive a variety of important applications, such as, SQL, Stream/Graph Processing and Machine Learning. A variety of different frameworks for expressing and running analytics workloads exist [2, 4, 5, 10, 12, 13, 16, 17]. These frameworks typically represent the workloads as DAGs of tasks that an underlying execution platform [11, 18, 21] runs across a cluster.

Despite their low-level differences, existing frameworks share the common attribute that they are intrinsically *compute-centric* (Figure 1). A typical framework starts by encoding exactly what computation is to be performed by each task, along with how many tasks to execute the same functionality and inter-task dependencies.[1] Given this computational structure, the framework then determines how to provide and route needed data to each task; typically, output data is hash-partitioned, with as many hash buckets as tasks to send data to. Finally, when to run a task is also tied to the computation structure; often, a consumer task is scheduled only when a fraction of its producer tasks have finished.

This compute-centric design has fundamental limitations. First, the rigidity of determining/specifying the computation's structure *a priori* limits the ability to express complex analytics. For example, a task's logic cannot be easily updated at runtime; at a more basic level, only limited changes to DAG structure can be made at runtime.

Second, coupling data exchange to compute structure may lead to suboptimal execution plans and hurt job performance; a small number of imbalanced partitions can lead to a large skew in data processing among tasks.

[1]This information is either specified by the user (e.g. spouts and bolts in Storm [17]) or automatically generated (e.g. a query optimizer [5]).

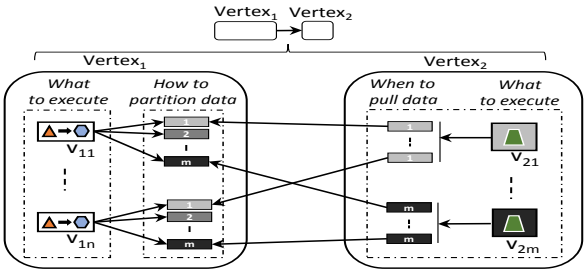

(a) Example of a job represented as a DAG with existing frameworks. $Vertex_1$ contains n tasks which perform the same two atomic operations; each task generates <Key,Value> pairs in m partitions. $Vertex_2$ has m tasks; every task aggregates a set of partitions with identical characteristics (e.g. same key range) and performs a single operation. This design is complex and hard to manage due to many assumptions regarding data semantics.

Vertex:
- Encodes set of operations to execute
- Decides task parallelism
- Decides when tasks can start executing

Edge:
- Encodes dependencies between vertices
- Defines number of partitions
- Defines routing of data between tasks
- Defines data lifetime

(b) Rigid requirements to materialize a job's execution. Note that they are often propagated to task level DAG.

Figure 1: Compute centric frameworks

Lastly, tying execution schedules to compute structure can lead to tasks unnecessarily hogging resources and hurt cluster efficiency.

In this paper we propose a new analytics framework F^2, which separates execution from data management, making compute and data equal first-class citizens. F^2 starts with the set of operations that are to be executed. Data is managed separately, and decisions to determine how data is partitioned or when it is to be processed are taken at runtime. In particular, the computation that processes the data, i.e., a "task", can have lose semantics and run any of the available operations on whatever data is ready, even across jobs.

This design offers several benefits. It improves flexibility in expressing analytics jobs by removing concerns regarding data partitioning, routing and what logic to specify early. It enables optimal task parallelism that is driven by the actual amount of generated data. It saves resources by triggering task execution only when partitions are ready. Finally, the generic tasks it allows can arbitrarily multiplex work, further improving efficiency and performance, and permitting pre-emption.

The primary goal of this paper is to make a case for decoupling compute and storage in big data analytics by highlighting both how it could be done and the benefits it offers relative to compute centricity. We start with an overview of our proposal.

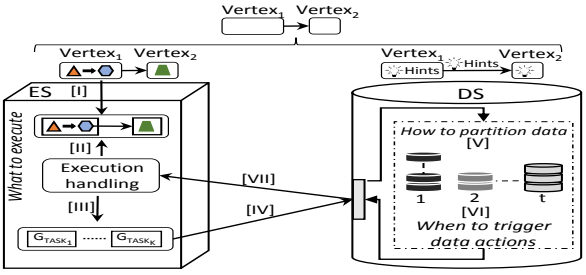

(a) F^2 divides a DAG's execution between ES and DS. This design provides significant simplifications w.r.t. to existing compute centric frameworks and enables more benefits.

ES	**Logical Vertex:** ▪ Encodes set of operations to execute	**Logical Edge:** ▪ Encodes dependencies between vertices
DS	**Logical Vertex:** ▪ Hints if it can execute on incomplete data	**Logical Edge:** ▪ Hints if special requirements how to do data partitioning

(b) Flexible requirements to materialize a job's execution.

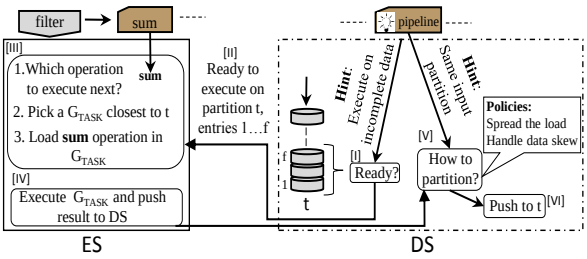

(c) Execution snapshot for a DAG where $filter$ (sum) are the current (next) operation to execute. F^2 encodes at DS that sum operation can execute on incomplete data. Due to this hint, when the partition t (with entries generated by $filter$ operation) has f entries, it notifies ES that partition t is ready for execution. ES checks in its logical DAG what is the next operation to be executed and finds sum. Subsequently, it picks a GTASK preferably on the same machine where partition t is located, loads the sum operation in the task and triggers its execution. Later on, when the corresponding output is generated, it is send to DS. Finally, DS collects statistics and has to decide how the new data is partitioned. To do that, it first analyzes if there are any hints w.r.t. this data. Otherwise it uses its own policies to spread the load or handle data skew. In this case, F^2 encodes at DS that the output of the sum operation can be pushed to the same partition used as its input (i.e. to t).

Figure 2: F^2

2 F^2 – High-level Design

F^2 uses existing programming models [15] to specify logical DAGs. Unlike existing approaches, it divides the data analytics functionality into two parts and translates DAG specifications accordingly (Figure 2).

An execution system (ES) is responsible for acquiring/releasing tasks from/to the underlying execution framework [11,18,21] and to execute DAGs' operations. It requires knowledge regarding which set of operations should execute together, and their order across vertices. To handle data management and trigger executions, F^2 provides a data system (DS). DS pushes data into *compute agnostic partitions* or CAPARTITIONS. These partitions are materialized in a single physical location.

F^2 works as shown in Figure 2a: At job submission time, it generates the logical DAG and encodes in ES the operations to execute and their dependencies (I). At runtime, given an input data, ES loads the next set of operations into a new or an existing task (II) and starts its execution (III). As soon as output is generated it is sent to DS (IV). DS stores the received data in CAPARTITIONS per DAG, using partitioning schemes driven by DS policies (e.g. to evenly spread load and handle data skew) and DAG specific hints (V). As soon as some data partition is ready (VI, e.g. all entries are received), DS informs the ES that new execution actions can happen given this input data (VII). This process repeats (from step II) until job completion. Figure 2c shows an execution snapshot to highlight F^2's internals.

This design provides significant simplifications (Figure 2b). For example, given a logical DAG, F^2 requires only information regarding what operations to execute and their dependencies. In contrast, a compute centric approach also requires information to materialize this execution at the task level (e.g. task parallelism, when tasks can be scheduled, etc.). F^2 hands off execution details to ES and takes task level actions driven by data semantics as directed by DS. Also, it does not require any data information (e.g. partitioning schemes, communication patterns, data lifetime policies, etc.) to be specified. However, if certain application requirements are to be enforced, F^2 uses hints to inform DS (e.g. a $OneToOne$ communication pattern may require all the data to go into the same partition).

F^2 enables many benefits. For example, as a consequence of a data driven design, tasks can be flexible enough to operate on any input data. F^2 uses GTASKS (Section 3.2) that multiplex execution across multiple partitions of data, vertices and different jobs.

Another benefit is that many operations can operate on incomplete data. F^2 design naturally enables support for pipelining operations due to the ability to structure the data in CAPARTITIONS and load operations across GTASKS. Figure 2c and Section 3 provide more details.

F^2 also enables data management at runtime. Due to CAPARTITIONS, F^2 can easily collect runtime statistics (e.g. the rate at which data is produced in a certain partition) and apply partitioning schemes to spread the load and handle data skew.

Next, we describe in detail how F^2 achieves these benefits, and many others.

3 F^2 – Benefits

We start with the benefits of decoupling execution from data management (§3.1).

3.1 Data Decoupling Benefits

Data-driven actions A common characteristic of existing data analytics frameworks is that computation cannot be executed until all the data is in place. For example, a consumer task has to retrieve all its assigned data before executing the set of operations. Because data aggregation can be very expensive, consumer tasks are often

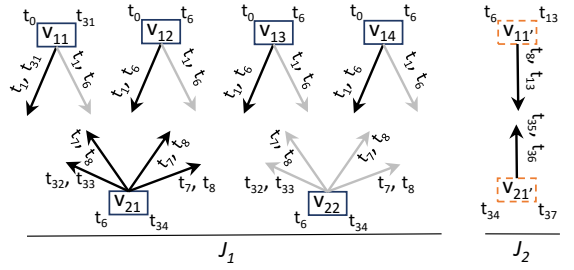

(a) Compute centric frameworks

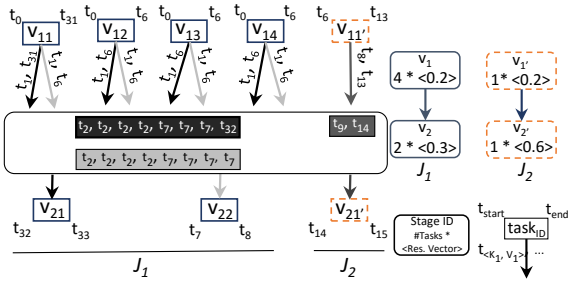

(b) F^2 (c) Two DAGs

Figure 3: Benefits of data-driven execution for two DAGs contending on a single bottleneck resource in a cluster of capacity 1. Jobs J_1 and Job J_2 (dashed lines) have two vertices each (3c). The legend shows: (1) the number of tasks in each vertex and their resource requirements; (2) the time when a task was launched/finished along with the times when new data entries have been made available for every data partition. Assuming that J_1 (J_2) starts at time 0 (6), (a)– (b) depict their execution: The average job completion time for a compute centric framework (a) is 33 time units and 21 time units for F^2 (b).

scheduled along with producer tasks to speedup data aggregation from tasks which are already completed. Despite its simplicity, this design choice does not necessarily help. Consider Figure 3a that computes the execution of the jobs described in 3c. J_1 assign all four tasks of v_1 at time t_0 which generate data entries organized in two partitions. Three of the tasks generate two entries per partition at t_1 and t_6 while the last one (v_{11}) generates the last entry at t_{31}. Observe that at time t_6 tasks of v_2 along with one task of $v_{1'}$ from J_2 are started. However, tasks of v_1 are not able to finish until task v_{11} completes. Thus, it hogs cluster resources which prevents task $v_{21'}$ from being scheduled. This leads to completion times of 35 and 31 time units for J_1 and J_2 respectively.

In contrast, F^2 pushes the data generated by each task of v_1 into two CAPARTITIONS as soon as they are generated. At J_2's submission time t_6 it assigns $v_{1'}$'s task ($v_{11'}$). When $v_{2'}$'s input is ready, it assigns task $v_{21'}$, leading to a completion time of 9 (3.44× better than before). Finally, when data is ready for v_2, its tasks are scheduled[2] and J_1 completes in 34 time units – or 1.53× better average job completion time.

F^2 outperforms existing approaches for two reasons: First, it triggers tasks' execution if and only if data is ready for them. Second, it overlaps data aggregation with data generation without hogging resources due to tasks doing aggregation (Figure 3a).

Note that, due to data aggregation in CAPARTITIONS, data is already in place when it is ready for processing. Subsequently, F^2 (ES) can collocate a task with its input data, reducing unnecessary network transfers[3]

The discussion above requires the entire input data to be in place before the corresponding computation can be applied. However, in practice many operations can pipeline their execution as new data arrives. Such tasks may require grouping and are commutative and associative (e.g. sum, min, $count$, $GroupBy$, $Join$, etc.). Pipelined execution increases tasks efficiency as tasks do not waste resources waiting for all data to accumulate. F^2 naturally enables pipelining (see Figure 2c for more details). For example, given several files containing <ProductId, SumSales> entries, and a partitioning scheme such that pairs with the same ProductId goes to the same partition, say we want to compute the total sum of sales for every ProductId. F^2 stores each CAPARTITION as a queue and performs the sum operation as follows: when the queue size increases above a threshold, it reads the records accumulated so far, performs the sum, pushes back the output, and deletes the previous entries. The new result along with new entries are further processed, till the data generation is complete.

Pipelining operations is not easily supported in existing systems. Some frameworks aim for simplicity, hence require that producer tasks sort their output before consumers can aggregate it, and to start applying operations only after all data is in place. Other approaches enable pipelined execution [2, 7, 9, 15, 17, 19, 21] but they are severely constrained. First, consumer tasks need to maintain local state regarding what data has been processed so far, and from which producer tasks. Second, it requires changes in the producer tasks to provide datastructures which allow easy separation of different subsets of data based on how it was processed by each consumer task. Finally, in case of failures neither consumer nor producer tasks are able to checkpoint what was already processed, leading to costly re-execution.

Data management at runtime Because F^2 treats data management as a standalone entity, it enables further optimizations at runtime. For example, because data entries of the same type are pushed to CAPARTITIONS as soon as they are generated, it is easy to track at runtime the rate at which partitions grow, due to which entries, and what are their sizes. This information enables F^2 to better *handle data skew* and *spread the load* across the machines. Specifically, if the operations are pipelined, parti-

[2]Although the input data for task v_{22} is ready at t_7, it can't be scheduled as task v_{11} can still produce further entries for it.

[3]This is not possible with existing approaches as tasks, reducers in particular, have to aggregate the data first; reducer placement is often random.

tion sizes remain small as tasks are executing along with data generation. Otherwise, DS keeps track of which keys are skewed and repartitions them into dedicated partitions. Similarly, it reacts to changes in load such that the partitions are spread equally across machines and their sizes are balanced.

With existing frameworks, developers have to write data partitioners [9, 15, 21] given knowledge regarding data distribution. However, accurate knowledge about data distribution is possible only for the initial data (e.g. input tables in a SQL database). Also, once a partitioning scheme is chosen, it cannot be changed at runtime. Finally, because partitions are physically present at multiple locations and tied with the tasks generating it, load can still be very imbalanced.

Better straggler mitigation Stragglers, tasks whose execution is much slower than other tasks in the same vertex, can significantly impact job performance. However, a systematic analysis of reasons for stragglers is challenging. To deal with this, data analytics frameworks launch speculative instances of the slow task which run in parallel. The instance which finishes first produces the final output. However, the speculative copy needs to execute on the entire data partition disregarding the progress already made by the straggler. This is because it is hard to checkpoint how far the execution went given the distribution of partition entries across multiple machines.

F^2's data partitioning can naturally delegate only certain data entries to the speculative instance. For example, if the slow task has made significant progress to process entries for keys $1...n$, then we only delegate entries with key range $n+1...m$ to one or more speculative copies. Thus, task execution is sped up *and* we minimize the amount of overlap in the data processed.

Runtime code simplification Existing frameworks often rely on runtime code generation to speed up execution. For example, [5,12] simplify the number of instructions required by a specific query plan to remove any overhead due to the query execution loading instructions that support broader-than-needed functionality. However these optimizations are mainly performed based on inspection of the query plan and initial data format rather than semantics of the generated data at runtime.

F^2 provides stronger opportunities for code simplification, by distilling the set of operators which are not required to be loaded for execution. For example, consider a query which performs an $OrderBy$ operation only on the values larger than 1000, where each value is produced through a join among two tables. At runtime, all the generated values in a partition are smaller than 1000. DS informs the ES that data is ready for execution along with various collected statistics, e.g. max value is below 1000. ES avoids launching unnecessary tasks that try to execute $OrderBy$ on this partition.

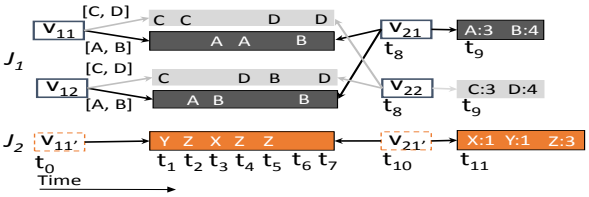

(a) Compute centric frameworks. J_1 generates data entries in two partitions between $t_1 \ldots t_7$. Finally, two consumer tasks aggregate the results. Similarly, J_2 generates data entries in a single partition. Note that despite $v_{1'}$ completes before v_1, tasks of v_2 are preferred for scheduling before $v_{2'}$ as they increase the cluster utilization the most. This leads to a completion time of 9 (11) time units for J_1 (J_2).

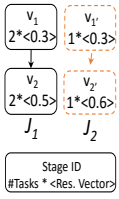

(b) Two $WordCount$ jobs DAGs

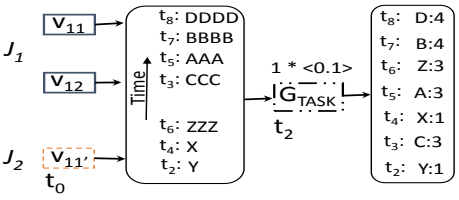

(c) F^2 using GTASKS. Whenever a partition is ready (all entries for a specific word are generated), *a single instance* of GTASK aggregates the results.

Figure 4: Benefits of GTASKS for two DAGs (b) which start at time 0, contending on a single bottleneck resource in a cluster of capacity 1. Using GTASKS (c), F^2 provides 1.42× better average job completion time than traditional approaches (a) using only one consumer task.

Changes to DAG structure Due to a rigid compute structure, existing analytics also have limited flexibility to adapt the existing DAG structure at runtime (e.g. eliminating vertices, change number of tasks). F^2 can provide further flexibility due to separation of DS and ES. For example, DS-collected statistics enable ES to migrate a $ShuffleJoin$ [6] operation from a consumer vertex to a $BroadcastJoin$ in the producer vertex.

Data-driven push based generic tasks Finally, F^2 simplifies DAGs' execution. Because DAG vertices do not hold information regarding vertex parallelism, which partition goes to which task or when to start requesting resources for a task, there is no tight coupling between any set of operations and a particular task. Consequently, a task in F^2 can operate on any partition of the data whenever it is ready to be executed. We call these *data-driven generic tasks* or GTASKS.

3.2 Benefits of GTASKS

GTASKS improve flexibility, efficiency and performance.

Runtime logic changes Due to a rigid compute structure, existing frameworks do not enable new execution logic to be specified at runtime. DAG operations can be modified only in the context of a new job instance. F^2 naturally enables this logic changes because, first, ES requires only the set of operations to execute (Figure 2b); the set can be modified at runtime. Second, GTASKS can load any functionality by design.

Improved efficiency, and performance To better understand these benefits, consider two jobs as described in Figure 4b. Figures 4a– 4c describe their execution using existing frameworks and F^2. Using GTASKS, F^2 provides $1.42\times$ better average job completion time than before. This is because F^2 enables efficient cluster use, which improves job performance. The efficiency gains come from GTASK properties.

First, GTASKS' multiplexing[4] ensures more efficient resource use than existing frameworks where: (i) task abstractions have well defined semantics (e.g. fixed input partition, fixed user logic, etc.) which can lead to suboptimal resource use at runtime, and (ii) a certain number of tasks are used, each of which is scheduled individually adding to scheduling overhead.

Second, GTASK enable F^2 to *scale* exactly with the amount of data and the rate at which the data is generated: given the number of partitions ready for execution, F^2 launches only the needed GTASKS.

Third, existing frameworks cannot interrupt a launched task before it finishes. This is because tasks have to aggregate data and keep state until their execution completes. However, if tasks can be easily interrupted, resources can be reallocated to critical jobs which can finish much faster. GTASK naturally facilitates *task interruption* as it does not maintain any local state and any other task can handle the same data as needed.

4 Discussion and Open Issues

Existing systems can be leveraged to implement F^2. For example, Kafka [1] can handle data partitions and maintain any required state. It provides reliability, high throughput and load distribution. GTASKS can be built atop HotTub JVM [8] which enables container reuse. Finally, we can use Tez's [15] DAG API and built-in communication with the underlying resource management framework [18].

Yet, there are several open issues. One key insight in our design is that various stages of execution in a job are triggered whenever data is ready. However, this is tightly related to how data is partitioned. We need a partitioning scheme where: (1) entries are grouped together such that data in each partition is ready at the earliest time; (2) each partition should have similar size; (3) we can change the number of partitions at runtime; (4) we have equal spread across machines; (5) partition placement minimizes network traffic w.r.t. to input sources and consumer tasks; (6) and we can repartition at runtime.

Another key challenge is to find the right datastructures to store data and any required state (e.g. data ready/not, how data is partitioned so far, statistics, active hints, etc). For example, partitions can be held in Kafka [1] queues. However, Kafka does not allow random reads/writes and deletions can be expensive.

Also, given that data is pushed out of the machine where it is generated, it should be stored reliably and able to be reproduced in case of failures. This process can significantly impact performance without proper consideration. For example, the amount of network/disk bandwidth contention due to bad replication can exceed the amount of resources used to regenerate data.

We must also consider the right abstractions for GTASKS. A generic task requires the ability to quickly react to changes in the execution as a consequence of data dynamics. This means bringing the task to a clean state after committing the result of its current execution, and loading a new set of operations. While HotTub [8] is a good start, the dynamics in execution changes can significantly impact Java Garbage Collector's performance and inherently slowdown task execution. An alternative is to use multithreaded processes [4, 12, 21]. However, the assumption here is the entire dataset fits in memory and tasks are short. Finally, an important challenge is how to configure the right amount of resources to be allocated to a GTASK given that its resource profile can change very often, e.g., based on current task logic.

5 Related Work

F^2 is not the first work to decouple data management from computation. [2,3] use Kafka [1] for high throughput communication and fault tolerance. Naiad [14] and StreamScope [20] come the closest to F^2 in the high-level principle of decoupling upstream and downstream vertices. The former simplifies task implementations, which can be agnostic to the degree of parallelism in a stage. The latter introduces abstractions to model computation on each vertex and to handle the data and communication aspects separately.

F^2 differs in at least three ways which make it more flexible and performant. First, it maps each vertex to an appropriate number of physical tasks at runtime, driven by the amount of input data and its arrival rate. Second, decisions regarding data partitioning and their allocation to tasks is driven by data. Third, tasks can perform arbitrary computation.

[15, 21] enable JVM reuse across dependent vertices in a job and HotTub [8] amortize the warm-up overhead across jobs. Other frameworks [4, 12, 21] avoid this overhead by launching long running containers and designating task execution to different threads. F^2 takes a step further and reduces this overhead by maximizing task efficiency. It multiplexes data partitions across running tasks and launches a new task only when the existing tasks are not able to keep up with the rate of partitions becoming ready for processing.

[4] across partitions, vertices and jobs; in our example GTASK executes on J_2's data whenever J_1's data is not ready

References

[1] Apache Kafka. https://kafka.apache.org/.

[2] Apache Samza. http://samza.apache.org.

[3] Kafka Streams. http://docs.confluent.io/3.0.0/streams/.

[4] Presto — Distributed SQL Query Engine for Big Data. prestodb.io.

[5] M. Armbrust, R. S. Xin, C. Lian, Y. Huai, D. Liu, J. K. Bradley, X. Meng, T. Kaftan, M. J. Franklin, A. Ghodsi, and M. Zaharia. Spark SQL: Relational data processing in Spark. In *SIGMOD*, 2015.

[6] S. Blanas, J. M. Patel, V. Ercegovac, J. Rao, E. J. Shekita, and Y. Tian. A comparison of join algorithms for log processing in mapreduce. In *SIGMOD*, 2010.

[7] T. Condie, N. Conway, P. Alvaro, and J. M. Hellerstein. MapReduce online. In *NSDI*, 2010.

[8] L. David, C. Adrian, S. Hailong, Z. Xin, G. Nikola, and Y. Ding. Dont get caught in the cold, warm-up your jvm: Understand and eliminate jvm warm-up overhead in data-parallel systems. In *OSDI*, 2016.

[9] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. In *OSDI*, 2004.

[10] J. E. Gonzalez, R. S. Xin, A. Dave, D. Crankshaw, M. J. Franklin, and I. Stoica. GraphX: Graph processing in a distributed dataflow framework. In *OSDI*, 2014.

[11] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. Joseph, R. Katz, S. Shenker, and I. Stoica. Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center. In *NSDI*, 2011.

[12] M. Kornacker, A. Behm, V. Bittorf, T. Bobrovytsky, C. Ching, A. Choi, J. Erickson, M. Grund, D. Hecht, M. Jacobs, I. Joshi, L. Kuff, D. Kumar, A. Leblang, N. Li, I. Pandis, H. Robinson, D. Rorke, S. Rus, J. Russell, D. Tsirogiannis, S. Wanderman-Milne, and M. Yoder. Impala: A modern, open-source SQL engine for Hadoop. In *CIDR*, 2015.

[13] X. Meng, J. K. Bradley, B. Yavuz, E. R. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. B. Tsai, M. Amde, S. Owen, D. Xin, R. Xin, M. J. Franklin, R. Zadeh, M. Zaharia, and A. Talwalkar. MLlib: Machine learning in Apache Spark. *CoRR*, abs/1505.06807, 2015.

[14] D. G. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, and M. Abadi. Naiad: A timely dataow system. In *SOSP*, 2013.

[15] B. Saha, H. Shah, S. Seth, G. Vijayaraghavan, A. Murthy, and C. Curino. Apache tez: A unifying framework for modeling and building data processing applications. In *SIGMOD*, 2015.

[16] A. Thusoo, R. Murthy, J. S. Sarma, Z. Shao, N. Jain, P. Chakka, S. Anthony, H. Liu, and N. Zhang. Hive – a petabyte scale data warehouse using Hadoop. In *ICDE*, 2010.

[17] A. Toshniwal, S. Taneja, A. Shukla, K. Ramasamy, J. M. Patel, S. Kulkarni, J. Jackson, K. Gade, M. Fu, J. Donham, N. Bhagat, S. Mittal, and D. Ryaboy. Storm@twitter. In *SIGMOD*, 2014.

[18] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, B. Saha, C. Curino, O. O'Malley, S. Radia, B. Reed, and E. Baldeschwieler. Apache Hadoop YARN: Yet another resource negotiator. In *SoCC*, 2013.

[19] A. Verma, B. Cho, N. Zea, I. Gupta, and R. H. Campbell. Breaking the MapReduce Stage Barrier. In *Cluster Computing*, 2013.

[20] L. Wei, Q. Zhengping, X. Junwei, Y. Sen, Z. Jingren, and Z. Lidong. Streamscope: Continuous reliable distributed processing of big data streams. In *NSDI*, 2016.

[21] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. Franklin, S. Shenker, and I. Stoica. Resilient Distributed Datasets: A fault-tolerant abstraction for in-memory cluster computing. In *NSDI*, 2012.